Subject: Re: User namespaces and keys
Posted by ebiederm on Wed, 23 Feb 2011 14:46:00 GMT
View Forum Message <> Reply to Message

"Serge E. Hallyn" <serge@hallyn.com> writes:

> Quoting David Howells (dhowells@redhat.com):
>>
>> I guess we need to look at how to mix keys and namespaces again.
>
>>From strictly kernel pov, at the moment, keys are strictly usable only
> by the user in your own user namespace.
>
> We may want to look at this again, but for now I think that would be a
> safe enough default. Later, we'll probably want the user creating a
> child_user_ns to allow his keys to be inherited by the child user_ns.
> Though, as I type that, it seems to me that that'll just become a
> maintenance pain, and it's just plain safer to have the user re-enter
> his keys, sharing them over a file if needed.
>
> I'm going to not consider the TPM at the moment :)
>
>> Possibly the trickiest problem with keys is how to upcall key construction to
>> /sbin/request-key when the keys may be of a different user namespace.
>
> Hm, jinkeys, yes.

Serge short term this is where I think we need to add a check or two so
that keys only work in the init_user_ns. When the rest of the details
are sorted out we can open up the use of keys some more.

As the first stage of converting the network stack we added patches
that turned off everything in non init namespaces. Those patches
were trivial and easy to review, and it made the conversion process
a lot easier. I suspect for keys and possibly security modules and
anything else that does is related we want to turn off by default.

Then once the core of user namespaces is safe to unshare without
privilege we can come back and get the more difficult bits.

Eric

_____
Containers mailing list
Containers@lists.linux-foundation.org
https://lists.linux-foundation.org/mailman/listinfo/containe rs